Co-tuning of a Hybrid Electronic-Optical Network for Reducing Energy Consumption in Embedded CMPs

Sandro Bartolini, Paolo Grani
Department of Information Engineering and Mathematical Sciences
via Roma 56, 53100, Siena, Italy
University of Siena
{bartolini, grani}@dii.unisi.it

ABSTRACT

Nanophotonic is a promising solution for on-chip interconnection due to its intrinsic low-latency and especially low-power features, desirable especially in future chip multiprocessors (CMPs) for rich *client* devices. In this paper we address the co-design of the parameters of a hybrid on-chip network featuring a traditional 2D mesh and a simple photonic helper ring aimed to improve performance and reduce energy consumption. As all the CMP traffic cannot be sustained in the considered simple optical interconnection without saturating the available bandwidth, and thus inducing performance and energy degradations, we identify the subset of coherency messages that are most worth to be accelerated through the low-energy optical path.

We investigate the management/arbitration strategies for the physically shared photonic path as they are crucial for reaching an effective exploitation of optical bandwidth according to their overhead and parallelism achieved in message transmission. Our results on multithreaded benchmarks, highlight that a careful selection of the most latency-critical messages to be routed on the photonic-path along with a *Multiple-Writers-Single-Reader* access scheme allows execution time and energy improvements up to 19% and 5%, respectively, for the 8-core setup and up to 16% and 13% for the 16-core configuration.

Furthermore, we show that the most aggressive ring access schemes allow the adoption of a four times slower electronic NoC that trades the achieved average speedup margin to obtain 70% overall energy savings, which is extremely important in energy constrained devices.

Categories and Subject Descriptors

C.1.2 [**Computer Systems Organization**]: Multiprocessors Interconnection architectures; B.4.3 [**Hardware**]: Interconnections Topology.

MES '13, June 23 - 24 2013, Tel-Aviv, Israel

General Terms

Performance, Design.

Keywords

On-chip optical interconnection, hybrid electro/optical network, low-power. tiled chip multiprocessors, token arbitration, high-end embedded architectures

1. INTRODUCTION

High-end embedded systems are evolving towards parallel architectures to keep the pace of Moore's Law prediction and the tiled paradigm is expected to enable design scalability [24]. Ideally, tiled CMP architectures comprise a number of identical tiles, each one having computational capabilities (e.g., a core), private and shared/distributed cache resources. Tiles are tied up with an on-chip interconnection network (NoC) that provides the communication and synchronization backbone for parallel application threads [12, 18, 20] on top of a directory-based cache coherence protocol.

This implies a central role of NoCs in future *client* devices - like smartphones, tablets, set top boxes and embedded PCs - from the performance and power consumption points of views. In particular, on-chip NoCs will have a crucial role in the future design of microprocessor systems and will need significant research investment [15] to overcome their limits in providing the expected performance within a reasonable power envelope. In particular, according to current technology and design roadmaps, on-chip electronic NoCs are expected to quickly account for about half of the overall chip power [9]. On-chip photonics is able to provide extremely low-latency and low-power communication

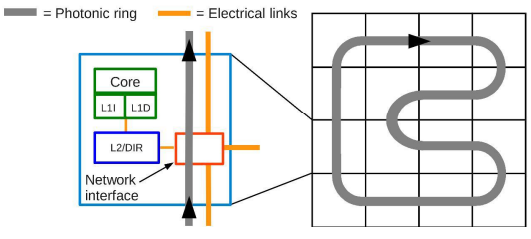

Figure 1: Overview of the considered tiled architecture for 16-core CMP setup. Every tile can transmit to and receive from the photonic ring.

but great care is needed in the design and management of photonic structures to translate these raw features into end-user perceived performance/power improvements within the usage pattern of a complex CMP running nowadays parallel applications.

As simple photonic topologies are likely to be the first to be integrated in near future embedded and *client* devices, in this paper we evaluate the maximum potential benefits that can be extracted from a photonic ring used in parallel (3D-stacked [19]) with a standard electronic NoC (eNoC). This kind of solutions allows avoiding articulated topologies (e.g., hundreds of waveguides) and/or complex management strategies (e.g., preliminary optical path setup) but cannot sustain all the CMP traffic in photonics without creating a bottleneck.

The selection of messages to be fruitfully forwarded onto the optical path is crucial to reach the right tradeoff between performance and energy consumption given the available parallelism and power capabilities. To access the shared photonic *medium* we analyzed three existing arbitration techniques. Each technique offers to message transmissions different bandwidth, arbitration overhead and communication parallelism. We analyzed their intrinsic performance and energy improvement capabilities. In particular, aiming at co-designing the features of the eNoC and the optical ring for maximum energy improvement, we show that *Multiple Writers Multiple Readers* and *Multiple Writers Single Reader* access schemes can either achieve significant execution time speedups (up to 20% and 28% for 8- and 16-core setups) or allow trading such speedup for slowing down four times the electronic NoC frequency while maintaining the original average performance. This coordinated design approach reduces up to 70% the overall energy consumption of the on-chip interconnection.

The remainder of the paper is organized as follows. In Section 2 we briefly introduce the technologies addressed by the paper. In Section 3, we describe the features of the considered architecture, while in Section 4 we present the evaluation methodology and discuss the achieved results. A survey of related works on photonic networks can be found in Section 5. Finally, Section 6 concludes.

2. ON-CHIP OPTICAL INTERCONNECTION

In this section we give some background about on-chip photonic interconnection technology. On-chip photonics is based on the effective possibility to integrate waveguides, lasers, modulators and photo-detectors into current process technology [23]. In particular, proposals exist adopting an on-chip optical network on a dedicated layer of a 3D-stacked chip architecture [22, 19]. The communication paradigm, diversely from electronic networks, is point-to-point by nature, as light packets cannot be managed in a store-and-forward fashion within the optical domain. Therefore, photonic links can be effectively used to directly connect communication endpoints. This feature determines one of the main benefits of this technology: power consumption of a transmission is, in practice, independent from the on-chip distance. For the same reason, communication latency is very fast because it is dominated by light propagation speed into silicon[1]. Passive [17] and active [23] optical switch architectures have been proposed for implementing network topologies with fixed and reconfigurable paths, respectively. Passive networks tend to require more optical modules than active ones to achieve similar connectivity. and active ones require preliminary optical path setup for micro-ring thermal tuning.

[1]About 66600 km/s $\rightarrow$ 15 ps/mm

Table 1: Parameters of the simulated architecture

Cores	8/16 cores, 2GHz
L1 caches	16kB(I)+16kB(D), 2-way, 1 cycle hit time
L2 cache	4 MB, 8-way, shared and distributed 8x512kB or 16x256kB banks. 3/12 cycles tag/tag+data
Directory	MOESI protocol, 8/16 slices, 3 cycles
eNoC (full-speed)	2D-Mesh, 2GHz, 5 cycles/hop
eNoC (low-power)	2D-Mesh, 0.5GHz, 5 cycles/hop
Photonic ring	3D, 30mm length, 8/16 I/O ports, 10GHz, 64 WDM, 460ps full round
Main memory	1GB, 300 cycles

3. ANALYZED ARCHITECTURE

Figure 1 exemplifies our baseline CMP architecture for 16-core setup where each core has private L1 caches (Instruction+Data) and a slice of L2 cache, which is shared and distributed between all cores. Directory information is distributed as well and is stored along each L2 cache slice. A traditional 2D electronic mesh allows communication between cores. Figure 1 shows also the photonic ring (3D-stacked), which is added to the baseline to obtain our considered enhanced architecture. We analyze 8- and 16-core architectures as representative of tiled CMPs close to the computational capabilities available today [1] or in the near future for high-end *client* devices like smartphones and tablets.

Table 2: Parameters of the electronic network

Parameter	**Fast eNoC**	**Slow eNoC**
Technology [nm]	32	32
Transistor type	NVT	HVT
Operating voltage [Volt]	1	0.8
Frequency [GHz]	2	0.5
Flit width [bit]	64	64
E_{link} [pJ/flit/mm]	8.98	4.95
$E_{switch-dyn}$ (5-port) [pJ/flit]	85.12	8.74
$E_{switch-dyn}$ (6-port) [pJ/flit]	118.79	15.87
$E_{switch-dyn}$ (7-port) [pJ/flit]	149.12	20.92
$P_{switch-static}$ (5-port) [mW]	33.60	5.17
$P_{switch-static}$ (6-port) [mW]	45.61	6.73
$P_{switch-static}$ (7-port) [mW]	58.08	8.35

We use dense wavelength division multiplexing (DWDM) [28] on the waveguide with 64 wavelengths modulated at 10GHz. Whenever a node has to send a message to another one, it access the ring and modulates the associated wavelength or wavelengths according to the photonic path and arbitration strategy. The photonic layer is thought to quickly route the most important part of the traffic generated by recent multithreaded applications and take advantage of the

Table 3: Parameters of the optical network

Frequency [GHz]	10
Flit width [bit]	64
$E_{dynamic}$ [pJ/bit]	0.41
$E_{logic-dynamic}$ [pJ/bit]	0.18
E_{serdes} [pJ/bit]	0.19
$P_{static-MWMR}$ (8-core) [mW]	183
$P_{static-MWMR}$ (16-core) [mW]	366
$P_{static-MWSR}$ (8-core) [mW]	73.8
$P_{static-MWSR}$ (16-core) [mW]	132
$P_{static-SWSR}$ (8-core) [mW]	73.6
$P_{static-SWSR}$ (16-core) [mW]	131

low-power capabilities of the photonic technology.

In particular we forward onto the photonic ring a carefully selected group of latency critical coherence messages of the MOESI protocol [5]. for improving performance. Specifically, these are the read-request messages (GETS), invalidations (INV) and invalidation acknowledgments (ACK). Furthermore, as in this work we aim at reducing energy consumption more than improving performance, we route onto the optical path also the actual data transmissions towards the requestors. These are not all data messages but only the most latency-critical ones. For example, writeback messages for dirty cache lines flow on the eNoC as execution time is less sensible to their transmission latency. The choice of including some data messages too aims at increasing the overall traffic on the optical channel and thus achieve higher energy saving due to the lower energy per bit of the photonic path than the electronic counterpart.

3.1 Considered Ring Arbitration Strategies

We analyzed three access strategies to the shared photonic ring: *Multiple Writers Multiple Readers* (MWMR), *Multiple Writers Single Reader* (MWSR) and *Single Writer Single Reader* (SWSR) [27, 28], according to the terminology introduced by Pan et al. [21]. Token ring is a well known protocol that guarantees starvation-free mutual exclusion on an ordered shared *medium* throughout the utilization of a special frame called *token*. In our setup the token is a light packet circulating on the ring. Token-acquisition (destructive read) grants the possessor permission to transmit on the *medium*. When a node in the NoC wants to transmit, it waits for the token, removes it from the ring (absorbing the light) and re-injects it again when it ends the transmission.

In a MWSR strategy, each receiver is associated to a predefined subset of all available colors (i.e., 8 lambdas for each receiver in the 8-core setup and 4 lambdas for each receiver in the 16-core configuration). A sender implicitly selects the destination node grabbing the correct token (the one associated to the node of interest) and transmitting on the destination-specific wavelengths (channel). In MWSR, any node can write to a given channel but only one node is able to read from it.

SWSR strategy has a single lambda (wavelength) statically assigned to a sender-receiver pair so it does not require arbitration at all (no token) and allows concurrent communications between all the cores (non-blocking network). The main issue of this design is its scalability. In a DWDM scheme with up to 64 lambdas on the sole waveguide available, SWSR can only manage up to eight network endpoints (e.g., cores). In fact, each endpoint (8) needs to have one lambda for communicating to each of the other endpoints (7), thus requiring $7 \times 8 = 56$ total lambdas. In a 16-core setup such a non-blocking network would require $16 \times 15 = 240$ lambdas, which is not feasible in our configuration with only one waveguide. The drawback of a static channel allocation between source-destination pairs is that a number of channels are idle if the traffic is not continuous and uniform between the endpoints, thus reducing the effective total network throughput when a specific path is needed. This is especially true in the cache coherent CMP scenario where cores need to transfer messages fast in certain moments but don't transmit in other time slots.

MWMR strategy instead can improve the ring utilization by sharing all lambdas among all network nodes. MWMR uses one token to access the ring and, before sending the message, it needs a receiver-selection phase to ask the destination to get ready for receiving. Once selected, all the available colors in the waveguide (64 in our setup) can be used for the transmission. The drawback of this design is that it needs a higher number of microrings because every node has to modulate and receive light in all the wavelengths. This implies to have more powerful laser sources to compensate the insertion-loss of the higher number of cascaded optical structures along the waveguide. Furthermore, even when idle, microrings still consume static power for thermal tuning.

The relative performance of such techniques are not obvious to predict. In fact going from MWMR to MWSR and SWSR the potential concurrency of different message transmissions can increases but optical parallelism of each one decreases, thus implying longer transmission time due to increased serialization effects.

4. RESULTS

In this section we introduce our evaluation methodology and test setup, and discuss the obtained results.

4.1 Methodology and Test Setup

Performance evaluations are obtained using the GEM5 simulator [8], in which we modeled both the all-electronic baseline and photonic enhanced tiled CMP architectures described in Section 3 and prepared a full-system setup running Linux 2.6.27 [13]. Table 1 summarizes some architectural details and parameters.

The total dynamic and static energy dissipated by all optical (lasers, microring resonators, modulators and photodetectors) and electronic modules (links and routers) have been taken into account for evaluating the overall consumption. Specifically eNoC energy parameters shown in Table 2 are derived by ORION 2.0 [16] for both full-speed (2GHz, Normal Voltage Threshold (NVT), 1.0V) and low-power (0.5GHz, High Voltage Threshold (HVT), 0.8V) configurations. In particular we need three kinds of electronic routers in our considered architectures (8- or 16-core): with 5 I/O ports (those on the mesh corners), with 6 I/O ports (those on the mesh borders) and with 7 I/O ports (those in the mesh middle for the 16-core setup). In fact each router needs to manage the mesh links plus three additional paths towards L1, L2 and directory modules. We choose ORION

2.0 instead of the more recent DSENT tool [26] for its native support for NVT and HVT CMOS technologies and, where possible (LVT), we validated the consumption parameters with DSENT tool.

ONoC energy parameters shown in Table 3 are in line with the state-of-the-art proposals. In particular for the serializer/deserializer (SERDES) circuits we conservatively scaled down the values in [11] from 65nm to 32nm technology. Furthermore we are using values from [30] for the energy consumption of the other optical devices (modulators, photo-detectors and microrings). Finally as a function of the architecture (8- or 16-core) and on the used access strategy (MWMR, MWSR and SWSR technique) we have calculated also all the static power dissipated by the lasers, thermal tuning of the microring resonators and the electronic interface circuitry to the photonic *medium*. In fact, topological implications of our setups are evaluated according to the detailed optical power estimation models developed internally thanks to the tight collaboration with optical technology groups within our national project on on-chip photonics. We verified that our results are in line with DSENT ones if our models are simplified to fit its ones and similar underlying photonic features are assumed.

Performance was evaluated for the PARSEC 2.1 benchmark suite [7, 4], a collection of heterogeneous parallel applications spanning different application domains and representative of diverse workloads that can be run also on advanced *client* devices (e.g., media processing, search and filtering, 3D and physics simulation). Benchmarks were instantiated with 8/16 threads, one per core and we used the medium input-set to maintain a reasonable simulation time while executing a good amount of instructions. As for performance metrics, we considered execution time of the entire parallel region of each benchmark as representative of the end-user perceived performance. Consumption metric is the overall network energy consumed by such parallel region. We were forced to neglect a few applications of the suite because of some problems in executing them in the considered simulation environment or due to scalability issues with the medium input size.

To have an idea of the behavior of the applications on our hybrid electro/optical NoC architecture, Table 4 shows the percentage of electronic and photonic traffic both for 8- and 16-core setups (number of packets and total number of transmitted bytes). For the 8-core configuration up to 87% of messages still remains on the eNoC and up to 84% for the 16-core one. Considering instead the transmitted bytes, these fractions are reduced to 75% and 70%, respectively, even if they are quite high.

This is due to two causes. First, only a subset of coherence messages are dynamically selected to be forwarded onto the photonic path. Secondly, in our setup the adoption of a *CLOSE/FAR* strategy is beneficial for overall performance even if it can reduce the traffic on the photonic path.

With this routing algorithm, a packet, if selected to go onto the photonic path, will be actually forwarded into it only if the number of eNoC hops to destination is greater than one. Otherwise (adjacent transmissions) eNoC is always preferred because is able to deliver the message maintaining an acceptable latency compared to an optical transmission and keeping at the same time the ring free for other communications. In fact, in our configuration, one eNoC hop requires five cycles at 2GHz (four for switch and one for link traversal) versus a minimum of three cycles for the optical path[2] (one for average token wait, one for optical network interface traversal and one for transmission management). Moreover this is the best case as the token could be far from the requestor (one additional cycle for its arrival) and that the photonic *medium* could be busy at the request time (causing additional waiting time). Hence, on average, one hop transmissions in optical can experience even higher latency than in the eNoC.

4.2 Result Discussion

Figure 2a and Figure 2b show the execution time and energy results in the 8-core setup and point out that the narrow bandwidth strategy (SWSR) is not the best choice as it induces an average 30% degradation both in time and energy results. Even if it lacks token arbitration overhead, its low-parallelism (1-bit per transmission) implies long serialization effects and thus message latency. This, in turn, slows down execution and therefore static energy, both electronic and optical, increases significantly.

Conversely, MWSR strategy delivers almost the same performance as the widest bandwidth MWMR strategy (21% vs 19% average improvement over the baseline) for the 8-core setup with a 5% improvement in energy consumption too. MWMR higher static energy consumption (see Section 3.1) erodes the dynamic energy advantage over the eNoC and scores 3% worse than the baseline. MWSR requires far less microring resonators as it splits them for receiving wavelengths among destinations (only eight photo-detectors for each receiver in the 8-core setup) thus obtaining a 7% energy improvement compared to the electronic baseline.

For the 16-core configuration, results for SWSR access strategy are missing from Figure 2 as it cannot be implemented within the 64 DWDM degree in only one ring as highlighted in Section 3.1. Figure 2c and Figure 2d show execution time and energy results, respectively, for the 16-core setup. MWMR access scheme is the best performing in execution time (an average 27% improvement over the baseline) while MWSR scores an average 18% improvement. However, MWSR scheme allows saving more energy than MWMR (13% vs 7%) due to its lower intrinsic static energy consumption.

In general, it is interesting to highlight that some benchmarks are able to take more advantage from the ultra low-latency optical path. For instance, *canneal* reaches more than 30% and 40% speedup for 8- and 16-core configurations. This is due to the high degree of fine-grain sharing and synchronization exposed by this benchmark which benefits from faster interaction between threads in invalidation and re-load of data among the average high number of sharers.

Table 4: Electronic and photonic traffic distribution

Simulation environment	E-traffic (packets/bytes)	O-traffic (packets/bytes)
8-core	87% / 75%	13% / 25%
16-core	84% / 70%	16% / 30%

We now want to show that the speedup obtained by the two wider bandwidth strategies (MWMR and MWSR), can

[2]MWMR, 64-bit coherence control message, ring free

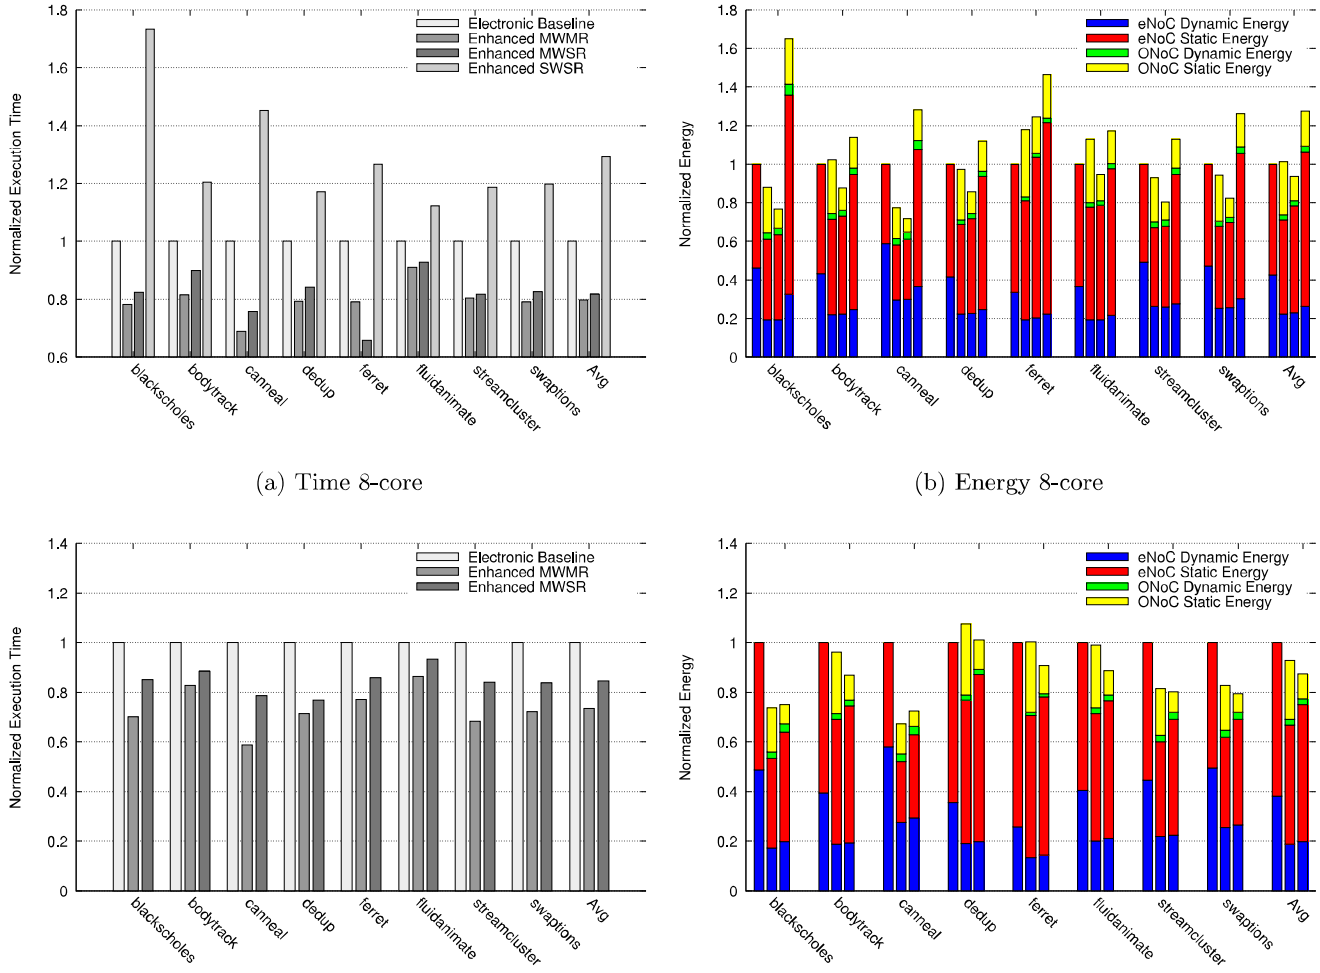

(a) Time 8-core

(b) Energy 8-core

(c) Time 16-core

(d) Energy 16-core

Figure 2: Comparison of the access management strategies on 8-core ((a) and (b)) and 16-core ((c) and (d)) configurations. In (b) and (d) bars are the same as in (a) and (c).

also be traded for a significant reduction of energy consumption, which is extremely interesting for the considered high-end *client* devices. through a careful co-design of the parameters of the electrical and optical networks. To get this, we looked for the slowest and more power efficient eNoC that could still match the baseline execution time performance. We identified that a four times slower electronic network, with lower-leakage transistor technology and assisted by the selected optical interconnection, could obtain the required performance-energy tradeoff. Figure 3 shows that this configuration is able to match within 3% the baseline execution time for both 8- and 16-core setups with both enhanced topologies (MWMR and MWSR). Some benchmarks, depending on their specific access pattern to shared and private data, show some slowdown or speedup. For example, *blackscholes* experiences a 28% slowdown for MWMR scheme on the 16-core architecture mainly due to his relatively small shared workload and higher fraction of access to private data with replacements. In practice, a higher fraction of latency-critical messages flows through the slower eNoC than in other benchmarks. Conversely *ferret* has a quite big shared working set and a high number of synchronization points in the pipelined algorithm. These are features that can induce significant speedup if the selected message set is accelerated through the photonic ring.

The frequency decrease in eNoC routers permits to scale down almost 10x their dynamic energy consumption and about 7x their static one, as shown in Table 2.

We finally want to remark that this electronic network slowdown, which can appear as a "*free meal*", is successful only if applied to the selected photonic enhanced architecture, which is able to offload from the eNoC the most latency-critical messages and deliver them quickly and with low-power towards destination. Otherwise, as shown in Figure 4 the low-power and 4 times slower eNoC alone induces more than 2x and up to 3x average slowdown for 8- and 16-core setup, respectively, despite being able to score around 70% energy reduction in both configurations. The results shown by Figure 3a and Figure 3b for the 8-core case confirm that the co-designed hybrid electro/optical network reaches

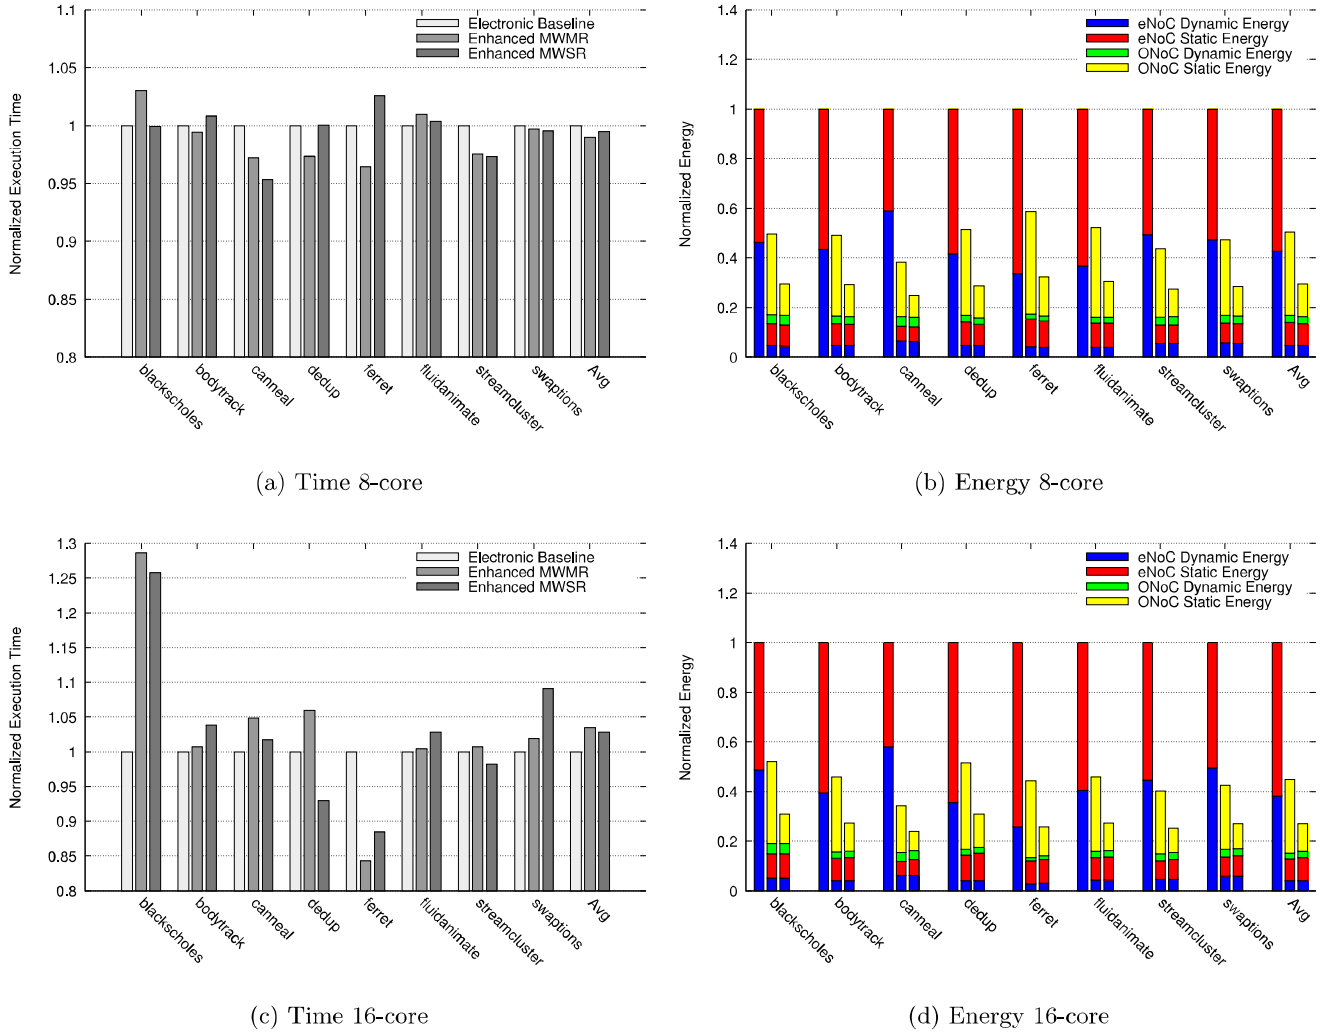

Figure 3: Comparison of the access management strategies on the photonic enhanced low-power architecture, with full-speed, all-electronic baseline on 8-core ((a) and (b)) and 16-core ((c) and (d)) configurations. In (b) and (d) the bars are the same as in (a) and (c).

50% energy average improvement for the MWMR access strategy and up to 70% for the MWSR one, which allows concurrent transmissions towards all cores but with 1/8 parallelism. For the 16-core architecture, energy results are in line with the 8-core ones. MWMR scores an average 56% energy improvement while MWSR access scheme reaches 72% energy reduction.

It is interesting to highlight that energy results are quite stable across the benchmarks and static quotes obviously can experience slight differences also according to the execution time deviation from the baseline performance.

5. RELATED WORK

Over the last years, there has been a growing interest in photonic interconnection (ONoC) as a means to alleviate the bandwidth and power consumption shortcomings of electronic NoC. Many photonic-based networks have been studied and several works have presented comparisons between on-chip electronic and photonic interconnects [25, 14, 3]. One attractive feature of photonic interconnection is the ability to operate in dense wavelength division multiplexing (DWDM), a technique which permits to transmit a high number (e.g., 64) of independent wavelengths in the same waveguide, improving the maximum potential bandwidth. Some recent works have proposed complete photonic NoCs which rely on DWDM capability. These solutions span from simple topologies that behave like crossbar based photonic NoC for on-chip communication like those presented by Vantrease et al. [28], Pan et al. [21], Pasricha et al [22], Bahirat et al [2] and Xu et al. [29], to more articulated ones like Petracca et al. [23] and Shacham et al. [25] that require or combine different transmission technologies. Most of such works address high-performance CMPs with a big number of cores and for this reason the required optical resources are quite high and the network are indeed quite complex. Instead of proposing new topologies, our work is focused on augmenting a standard CMP, endowed with a 2D mesh eNoC, with a single photonic ring for data transmis-

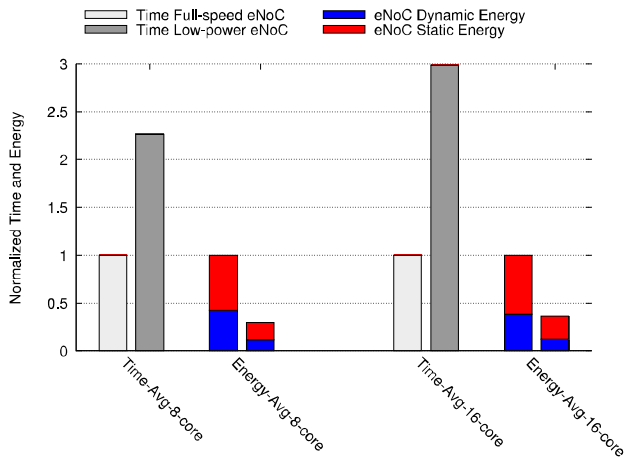

Figure 4: Comparison of time and energy consumption results for all-electronic (without optical path) mesh architecture with full-speed and low-power eNoC for 8- and 16-core configurations.

sion and then exploiting the best features of nanophotonics while avoiding the negative effects of its shortcomings, with the global objective of decreasing the overall energy consumption of the hybrid on-chip electro/optical interconnection.

Regarding the architectures with simpler topologies, Pan et al. [21] proposed a cluster of nodes in which inter-cluster communication is done using similar crossbar architecture. Other researchers [28, 21, 29] present very large architectures with up to thousand nodes and a some hundreds of optical waveguides. Furthermore Xu et al. [29] is more focused on coherence issues proposing a composite cache coherence protocol that benefits from direct cache-to-cache accesses as in snoopy protocol and small amount of cache probing as in directory-based protocols. Pasricha et al [22] proposed a 3D-stacked optical simple architecture that combines multiple photonic rings on various photonic layers with a 3D mesh NoC in active layers. Unlike the works discussed above, in this paper we propose to extend a traditional all-electronic mesh NoC with only one photonic ring. The simple ring waveguide reduces the overall network complexity compared to the cited works. Bahirat et al [2] also proposed an electronic mesh augmented by a number of photonic rings, on different 3D-stacked layers, that are used to facilitate global on-chip communication between distant processors and memory cores in high-performance CMPs. In our work, instead, the ring is used to provide a low-latency and high-bandwidth path to a specifically selected set of most latency-critical control messages of a MOESI coherency protocol for a near-future embedded tiled architecture and we maintain full eNoC connectivity (no clusterization) between the cores.

Within proposals with more complex topologies, some [23, 25] employ a preliminary optical path setup through a helper electrical networks (active networks). This can potentially limit the latency and energy advantages of photonics depending on the lifespan of optical paths vs setup overhead and due to possible waiting time for contended sub-paths. Instead. we focus on a simple photonic structure (ring) as, if it is properly managed, it can potentially deliver significant latency and energy improvements without needing big investments in complex and/or articulated photonic structures.

Other proposals implement passive wavelength routing [10], which relies on statically configured optical switches and WDM technique to obtain conflict-free full connectivity between all network endpoints. All paths connecting every source-destination pair can be used concurrently but the degree of parallelism of each one is limited (e.g., one bit). In our proposal we aim at employing simpler optical topologies and we found that, especially for supporting coherency protocol evolution, higher parallelism in each transmission is more important for avoiding slowdowns, even if at a certain arbitration cost. than providing contention-free network with limited parallelism.

From another perspective, existing works and techniques for effectively exploiting different electronic networks are not suitable in this hybrid electro/optical domain as typically faster (lower latency) eNoCs are the more power hungry than the slower ones. Using a photonic helper network changes the rules of the game as the optical path has ultra low-latency and low dynamic energy consumption. But beyond a certain traffic level latency, and consequently energy, benefits fade. Therefore, identifying the sweet spot between the network parameters and the traffic that allows achieving the best results is delicate and it is something different from any other approach working with only electronic networks. To the best of our knowledge, this is the first work that aims at co-designing the parameters of a hybrid electro/optical network with limited optical resources for pursuing maximum energy saving in high-end embedded CMPs.

6. CONCLUSIONS

Photonic-electronic integration expects to change the landscape of future computing and networking [6], especially for pursuing energy efficiency. In this paper we presented a study of co-optimization of the parameters of a hybrid electro/optical network for tiled CMP architectures enabling near future *client* devices, with the objective of identifying low-energy solutions. We analyzed the cross interactions between well known access strategies (MWMR, MWSR and SWSR) for accessing a single shared ring structure, working in parallel with the electronic NoC, and the design parameters of the eNoC itself. The investigation has been performed in the context of recent multithreaded applications running on a full-system simulator. We demonstrated that a careful selection of latency-critical and energy-sensitive messages of the adopted coherence protocol can be routed to the photonic ring to obtain 19% performance and 5% power advantages on the 8-core configuration. On 16-core one these values become 16% and 13%, respectively. Then, we highlighted that such specific traffic unloaded from the eNoC could allow to relax its requirements. This way, we showed that a four times slower eNoC could match the original baseline performance while enabling about 70% reduction in the energy consumption of the overall on-chip network.

7. ACKNOWLEDGMENT

This work was partially supported by IT FIRB Photonica (RBFR08LE6V) and IT PRIN 2008 (200855LRP2) projects.

8. REFERENCES

[1] Samsung exynos 5 octa, 2013.

[2] S. Bahirat and S. Pasricha. Exploring hybrid photonic networks-on-chip foremerging chip multiprocessors. In *International conference on Hardware/software codesign and system synthesis*, 2009.

[3] K. Bancrjcc, A. Mchrotra, A. Sangiovanni-Vincentclli, and C. Hu. On thermal effects in deep sub-micron vlsi interconnects. In *Design Automation Conference*, pages 885–891, New York, NY, USA, 1999. ACM.

[4] N. Barrow-Williams, C. Fensch, and S. Moore. A communication characterisation of splash-2 and parsec. In *IISWC*, 2009.

[5] S. Bartolini and P. Grani. A simple on-chip optical interconnection for improving performance of coherency traffic in cmps. In *Proceedings of the Digital System Design, DSD*, Cesme, Turkey, Sept. 2012.

[6] S. Ben Yoo. Impact of photonic-electronic integration in networking and computing systems. In *Opto-Electronics and Communications Conference (OECC), 2012 17th*, pages 293 –294. july 2012.

[7] C. Bienia, S. Kumar, J. P. Singh, and K. Li. The parsec benchmark suite: Characterization and architectural implication. In *PACT*, Oct. 2008.

[8] N. L. Binkert, R. G. Dreslinski, L. R. Hsu, K. T. Lim, A. G. Saidi, and S. K. Reinhardt. The m5 simulator: Modeling networked systems. *IEEE Micro*, 26, 2006.

[9] S. Borkar and A. A. Chien. The future of microprocessors. *Commun. ACM*, 54(5):67–77, May 2011.

[10] M. Brière, B. Girodias, Y. Bouchebaba, G. Nicolescu, F. Mieyeville, F. Gaffiot, and I. O'Connor. System level assessment of an optical noc in an mpsoc platform. In *Proceedings of the conference on Design, automation and test in Europe*, DATE '07, pages 1084–1089, San Jose, CA. USA, 2007. EDA Consortium.

[11] R. Dobkin, M. Moyal, A. Kolodny, and R. Ginosar. Asynchronous current mode serial communication. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 18(7):1107 –1117, july 2010.

[12] H. Franke, M. Kirkwood, and R. Russell. Fuss, futexes and furwocks: Fast userlevel locking in linux. In *Ottawa Linux Summit*, 2002.

[13] M. Gebhart, J. Hestness, E. Fatehi, P. Gratz. and S. W. Keckler. Running parsec 2.1 on m5. Technical report, The University of Texas at Austin, Department of Computer Science, October 2009.

[14] J. W. Goodman, F. J. Leonberger, S. Y. Kung, and R. A. Athale. Optical interconnections for vlsi systems. In *IEEE*, volume 72, pages 850–866, Jul. 1984.

[15] ITRS. International technology roadmap for semiconductors - interconnection. *ITRS*, 1, 2011.

[16] A. Kahng, B. Li, L.-S. Peh, and K. Samadi. Orion 2.0: A power-area simulator for interconnection networks. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 20(1):191 –196, jan. 2012.

[17] S. Koohi, A. Shafaei, and S. Hessabi. An optical wavelength switching architecture for a high-performance low-power photonic noc. In *Workshops of Int.l Conference on Advanced Information Networking and Applications*, 2011.

[18] Z. Li, J. Wu. L. Shang, A. Mickelson, M. Vachharajani, D. Filipovic, W. Park, and Y. Sun. A high-performance low-power nanophotonic on-chip network. In *Int. Symp. Low Power Electronics and Design*, Aug. 2009.

[19] G. H. Loh. 3d-stacked memory architectures for multi-core processors. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*, pages 453–464, Washington, DC, USA, 2008. IEEE Computer Society.

[20] I. O'Connor and F. Gaffiot. On-chip optical interconnect for low-power. In E. Macii, editor, *Ultra-Low Power Electronics and Design*, Kluwer, Dordrecht, 2004.

[21] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary. Firefly: Illuminating future network-on-chip with nanophotonics. In *Proceedings of the Int.l Symposium on Computer Architecture*, 2009.

[22] S. Pasricha and S. Bahirat. Opal: a multi-layer hybrid photonic noc for 3d ics. In *ASPDAC*, Jan. 2011.

[23] M. Petracca, B. G. Lee, K. Bergman, and L. P. Carloni. Design exploration of optical interconnection networks for chip multiprocessor. In *IEEE HOTI*, Washington, DC, USA, 2008.

[24] A. Ros, M. E. Acacio, and J. M. Garc. Scalable directory organization for tiled cmp architectures. *ICCD*, 2008.

[25] A. Shacham, K. Bergman, S. Member, and L. P. Carloni. Photonic networks-onchip for future generations of chip multiprocessors. *IEEE Trans. Computing*, page 1260, 2008.

[26] C. Sun, C.-H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic. Dsent - a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling. In *Proceedings of the 2012 IEEE/ACM Sixth International Symposium on Networks-on-Chip*, NOCS '12, pages 201–210, Washington, DC, USA, 2012. IEEE Computer Society.

[27] D. Vantrease, N. Binkert, R. Schreiber, and M. H. Lipasti. Light speed arbitration and flow control for nanophotonic interconnects. In *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO 42, pages 304–315, New York, NY, USA, 2009. ACM.

[28] D. Vantrease, R. Schreiber, M. Monchiero, M. Mclaren, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn. Corona: System implications of emerging nanophotonic technology. In *ISCA*, June 2008.

[29] Y. Xu, Y. Du, Y. Zhang, and J. Yang. A composite and scalable cache coherence protocol for large scale cmps. In *Proceedings of the international conference on Supercomputing*, ICS '11, pages 285–294, New York, NY, USA, 2011. ACM.

[30] X. Zheng and al. Ultra-efficient 10gb/s hybrid integrated silicon photonic transmitter and receiver. *Opt. Express*, 19(6):5172–5186, Mar 2011.